

The angular momentum of cold dark matter haloes with and without baryons

Philip Bett,^{1,2*} Vincent Eke,¹ Carlos S. Frenk,¹ Adrian Jenkins,¹
and Takashi Okamoto^{1,3}

¹*Institute for Computational Cosmology, University of Durham, South Road, Durham, DH1 3LE, UK*

²*Argelander-Institut für Astronomie, Universität Bonn, Auf dem Hügel 71, D-53121 Bonn, Germany*

³*Center for Computational Sciences, University of Tsukuba 1-1-1, Tennodai, Tsukuba, Ibaraki 305-8577, Japan*

18 January 2010

ABSTRACT

We investigate the magnitude and internal alignment of the angular momentum of cold dark matter haloes in simulations with and without baryons. We analyse the cumulative angular momentum profiles of hundreds of thousands of well resolved haloes in the Millennium simulation of Springel et al. and in a smaller, but higher resolution, simulation, in total spanning 5 orders of magnitude in mass. For haloes of a given mass, the median specific angular momentum increases with radius as $j(\leq r) \propto r$. The direction of the vector varies considerably with radius: the median angle between the inner ($\lesssim 0.25R_{\text{vir}}$) and total ($\leq R_{\text{vir}}$) angular momentum vectors is about 25° . To investigate how baryons affect halo spin, we use another high-resolution simulation, which includes gas cooling, star formation and feedback. This simulation produces a sample of galaxies with a realistic distribution of disc-to-total ratios, D/T : two thirds of the galaxies have $D/T > 0.5$ in the B-band. The formation of the galaxy spins up the dark matter within $0.1R_{\text{vir}}$ such that the specific halo angular momentum increases by ≈ 50 per cent in the median. The dark matter angular momentum becomes better aligned, but there remains a broad distribution of (mis-)alignments between the halo and the central galaxy, with a median angle between their angular momenta of $\sim 30^\circ$. Galaxies have a range of orientations relative to the shape of the halo: half of them have their minor axes misaligned by more than 45° , although only about 10 per cent of the galaxies lie within 30° of the plane perpendicular to the major axis of their halo. Finally, we align a sample of haloes according to the orientation of their galaxies and stack the projected mass distributions. Although the individual haloes are significantly aspherical, galaxy–halo misalignments produce a stacked mass distribution that cannot be distinguished from circular. If the lack of alignment found in our simulations is realistic, it will be extremely difficult for weak lensing studies to measure the ellipticity of cold dark matter haloes using this technique.

Key words: cosmology: dark matter – galaxies: haloes – methods: N -body simulations

1 INTRODUCTION

The formation of structures in the universe is often studied in the context of the two stage model of White & Rees (1978), within the Λ cold dark matter (Λ CDM) cosmology (White & Frenk 1991): primordial mass density perturbations (dominated by cold dark matter or CDM) collapse under gravity, hierarchically forming ever larger haloes, while galaxies form from the baryons that cool and collapse within them. The first part of the problem – determining the structure and evolution of dark matter haloes – can be tackled realistically and reliably using N -body simulations. Starting from the early simulation work of Davis et al. (1985) and Frenk et al. (1985,

1988), there is now a vast amount of literature on the subject, covering both detailed studies of a few individual objects at very high resolution (e.g. Navarro, Frenk & White 1996, Okamoto & Habe 1999, Knebe et al. 2000, Bullock et al. 2001, Springel et al. 2005, Diemand, Kuhlen & Madau 2007, Springel et al. 2008), as well as statistical studies of large samples of less well resolved haloes (e.g. Bett et al. 2007, (hereafter B07), Hahn et al. 2007, Macciò et al. 2007). Some of these studies have focused specifically on the angular momentum structure of haloes, which is the main topic of this paper (e.g. Davis et al. 1985, Barnes & Efstathiou 1987, Warren et al. 1992, Cole & Lacey 1996, Bullock et al. 2001, Porciani et al. 2002a,b, Chen & Jing 2002, Vitvitska et al. 2002, Kasun & Evrard 2005, Bailin & Steinmetz 2005, Avila-Reese et al. 2005,

* Email: p.e.bett@dunelm.org.uk

Shaw et al. 2006, Allgood et al. 2006, B07, Macciò et al. 2007, 2008).

The dissipative baryonic processes that produce the visible galaxy are considerably more complex, and are correspondingly less well understood, than the purely gravitational processes that make the dark matter halo. Nevertheless, recent N -body/gasdynamics simulations, which include various forms of feedback between the cooling gas and the forming galaxy, are beginning to produce fairly (although not entirely) realistic disc galaxies from Λ CDM initial conditions. Recent examples of such work include Abadi et al. (2003a,b), Sommer-Larsen, Götz & Portinari (2003), Robertson et al. (2004), Governato et al. (2004, 2007), Okamoto et al. (2005), Brook et al. (2006), Tasker & Bryan (2006), Gustafsson, Fairbairn & Sommer-Larsen (2006), Heller, Shlosman & Athanassoula (2007), Ceverino & Klypin (2008), Croft et al. (2009), Robertson & Kravtsov (2008), Gibson et al. (2009), Scannapieco et al. (2009), Romano-Díaz et al. (2009); see also Okamoto (2008) and Mayer, Governato & Kaufmann (2008) for recent reviews.

Baryonic processes modify the internal structure of the halo in which they take place, including basic properties such as their angular momentum and shape profiles. These are the quantities that we study in this paper. We first exploit the enormous statistical power of the Millennium Simulation (Springel et al. 2005) to investigate pure dark matter haloes with a large range of masses, from galactic to cluster haloes. We then compare the shape and angular momentum properties of galactic size haloes with those of haloes in a smaller, higher resolution simulation that includes baryonic physics. In addition to the intrinsic importance of these fundamental properties, we are interested in this problem for two practical reasons. Firstly, even the best gasdynamics simulations to date appear to suffer from the “angular momentum problem” first highlighted by Navarro & Benz (1991). This is a statement of the fact that in simulations, excessive transfer of angular momentum from the cooling gas to the halo leads to discs that are much smaller than observed. On the other hand, calculations of galaxy formation using semianalytic techniques, which assume that the baryons initially have the same angular momentum distribution as the dark matter, and that this is conserved as the baryons collapse produce better agreement with observed disc sizes (see, e.g. Fall & Efstathiou 1980; Mo et al. 1998; Cole et al. 2000; de Jong & Lacey 2000; Dutton et al. 2007; González et al. 2009).

The second reason why we are interested in halo shapes and angular momentum is the prospect of measuring the shapes of real galactic haloes using weak gravitational lensing (see Hoekstra & Jain 2008, for a recent review). In principle, the shapes of dark matter haloes provide an important test of the Λ CDM model. Simulations that neglect baryon effects have demonstrated that cold dark matter haloes are generically triaxial (Frenk et al. 1988; Dubinski & Carlberg 1991; Warren et al. 1992; Jing & Suto 2002; Allgood et al. 2006; Shaw et al. 2006; Bett et al. 2007; Hayashi et al. 2007). On the other hand, some alternative gravity theories such as TeVeS/MOND (Bekenstein 2004; Mörklock & Turner 2001) predict the gravitational potential at large distances from galaxies to be spherical. Since the gravitational lensing signal from an individual object is very weak, lensing studies are normally performed on stacked images around large samples of galaxies, orientated according to the shape of the galaxy. Intrinsic to this method is the assumption that the shape of the galaxy and the shape of the halo are tightly correlated with each other. It is not clear, however, that cold dark matter theory predicts this.

Recent attempts to measure halo shapes from weak lensing

data have been inconclusive. Hoekstra et al. (2004), using the Red Sequence Cluster Survey (RCS), claimed to detect a definite halo ellipticity; Mandelbaum et al. (2006), using data from the Sloan Digital Sky Survey (SDSS), were unable to obtain a definite detection of halo ellipticity, but their data suggested different galaxy–halo alignment distributions for spiral and elliptical galaxies. Parker et al. (2007), using the CFHT Legacy Survey, again detected an aspherical halo shape, but at relatively low significance. Evans & Bridle (2008), using a similar method but applied to cluster haloes, claimed to rule out spherical configurations.

There are, of course, other methods for measuring halo shapes, including using cluster X-ray data in conjunction with the Sunyaev–Zel’dovich effect (e.g. Reblinsky 2000; Zaroubi et al. 2001; Lee & Suto 2004; Sereno et al. 2006), or examining the distribution of satellite galaxies within a halo (e.g. Plionis et al. 1991; Basilakos et al. 2000; Wang et al. 2008). These are not considered in this paper.

Several studies have investigated the influence of baryonic processes on the shapes and angular momentum of haloes. The case of non-radiative gas was considered by van den Bosch et al. (2002) and Sharma & Steinmetz (2005), who found a broad distribution of dark matter/gas orientations. Van den Bosch et al. (2002) also found significant misalignment between the inner and outer regions of the halo. Kazantzidis et al. (2004), Springel, White & Hernquist (2004), and Bailin et al. (2005) all showed that haloes become much more spherical in simulations with gas cooling and star formation, particularly in the central regions, but with a significant effect throughout the halo. This radial trend is the opposite of that found in simulated haloes without baryons: haloes simulated with just dark matter become *less* spherical (more prolate) towards their centres (Dubinski & Carlberg 1991; Hayashi et al. 2007). The constrained realisation simulations of Berentzen & Shlosman (2006) showed that it is the growth of the baryonic disc that reduces the halo prolateness, pushing it closer to spherical¹. More recently, Abadi et al. (2009) compared 13 haloes simulated with and without gas (including radiative cooling, but without star formation or feedback), and broadly confirmed these results.

Early investigations of the alignment between the *angular momentum* vectors of the halo and gas distributions were performed by van den Bosch et al. (2002), Yoshida et al. (2003), and Chen, Jing & Yoshikawa (2003), all of whom considered both cooling and non-radiative gas. They found that the cold gas develops a broader range of orientations relative to the dark matter than either the hot or non-radiative gas. Bailin et al. (2005) considered the alignment of simulated galaxies with their parent halo. They found that, although the dark matter at the virial radius remained essentially uncorrelated with the orientation of the galactic disc, the presence of the disc had caused the minor axis of the inner halo to align with the disc axis. Gustafsson, Fairbairn & Sommer-Larsen (2006) similarly examined the density, shape and orientation profiles of haloes in simulations with and without baryonic physics. They too found that the baryons made their haloes more spherical, with a wide range of angles between the dark matter and the galaxies. Van den Bosch et al. (2003) considered the effect of preheating of the intergalactic medium and found that it greatly increased the misalignment

¹ This behaviour is not universal however. In the single object simulated by Heller et al. (2007), the inner regions of the halo are consistently more spherical than the outer halo, regardless of whether baryons were included or not.

between the angular momentum of the (non-radiative) gas and the dark matter halo.

Croft et al. (2009) used a model that included black hole feedback in the star formation process, and found a broad distribution of angles between the orientations of the galaxies and their parent haloes, as well as between the shape and angular momentum axes of each halo component (gas, stars and dark matter). Most recently, Romano-Díaz et al. (2009) compared simulations of a single halo, with and without baryons and star formation, and found that the dark matter angular momentum was very well aligned to the stellar component throughout the halo, but the gaseous component was $\gtrsim 100^\circ$ out. However, the alignment of the angular momentum vectors in their halo varied dramatically as the halo evolved.

While these simulations tend to produce broadly consistent results, they cannot easily be compared because of the different type of baryonic processes that they included. An important limitation of these studies is the small samples of objects they were able to simulate, which preclude robust statistical statements. The one exception to this is the large simulation by Croft et al. (2009) which, however, had to be stopped at $z = 1$ due to computing time restrictions.

In the first part of this paper, we exploit the statistical power of the Millennium Simulation (Springel et al. 2005) to determine the properties of the angular momentum profiles of dark matter haloes of a large range of mass in a purely cold dark matter universe. In Section 2, we consider both the magnitude and orientation of the angular momentum vectors as functions of halo radius and mass. In Section 3, we include the effects of baryons and explore a pair of simulations of a smaller volume, one of which contains just dark matter and the other which contains dissipative baryons, star formation and feedback as well. We examine the angular momentum profiles of the haloes in detail, comparing the dark matter at different radii between the two simulations, and their orientation relative to the central galaxy. We also examine the physical reasons behind the changes induced by the baryons in the dark matter. Finally, we apply our results directly to the problem of weak gravitational lensing by computing the 2-D projected mass distributions (individual and stacked), when haloes are aligned according to their galaxy's orientation. We present our conclusions in Section 4.

2 DARK MATTER HALOES

We begin by examining the angular momentum structure of simulated dark matter haloes without baryons.

2.1 The simulations

We analyse two dark matter simulations of the large-scale structure of a Λ CDM universe in order to obtain very precise statistics of dark matter halo properties over a wide range of masses. The first of these, the *Millennium Simulation* (MS), has over 10 billion particles (2160^3) in a $500 h^{-1}\text{Mpc}$ cubic volume; it is described fully in Springel et al. (2005). The second simulation (which we will refer to as hMS) assumes the same cosmology, but has higher resolution and a smaller volume². Both simulations were carried out with a version of the GADGET-2 code (Springel 2005) that

² The hMS simulation was also analysed by Neto et al. (2007) and Gao et al. (2008).

Sims.	Ω_Λ	Ω_M	Ω_b	h	n	σ_8
MS & hMS	0.75	0.25	0.045	0.73	1.0	0.9
DMG & DMO	0.70	0.30	0.044	0.70	1.0	0.9

Table 1. Cosmological parameters (at $z = 0$) for the simulations used in this paper. The cosmological density parameters are defined as $\Omega_i := \rho_i/\rho_{\text{crit}}$, where the critical density $\rho_{\text{crit}} := 3H_0^2/(8\pi G)$, the equivalent mass density of the cosmological constant is $\rho_\Lambda := \Lambda c^2/(8\pi G)$, and $\Omega_{\text{tot}} = \Omega_\Lambda + \Omega_M = 1$. The Hubble constant at $z = 0$ is parametrised as $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$. The spectral index is given by n , and σ_8 is the linear theory mass variance in spheres of radius $8 h^{-1}\text{Mpc}$ at $z = 0$. (The DMO and DMG simulations are described in Section 3.)

Sim.	L_{box} $h^{-1}\text{Mpc}$	N_{part}	m_p $10^7 h^{-1}\text{M}_\odot$	η $h^{-1}\text{kpc}$
MS	500	10 077 696 000	86.07	5.0
hMS	100	729 000 000	9.52	2.4
DMO	35.33	3 397 215	1.93	0.7
DMG:	35.33			
–DM		3 397 215	1.65	0.7
–Gas		2 985 242	0.28	0.35
–Stars		1 668 836	0.06	0.35

Table 2. Parameters for the simulations used in this paper: box size, numbers and masses of particles, and gravitational softening η (see equation 4). For DMG and DMO, we analyse the high resolution region in the centre of the L_{box} cube; at $z = 0$, this region is approximately spherical with a diameter of about $12.5 h^{-1}\text{Mpc}$. Note also that in DMG, the number of gas and star particles, and their masses, vary over the course of the simulation according to the star formation algorithm. The values presented here are the numbers and median masses at $z = 0$.

was specially optimised for massively parallel computations and low memory consumption.

We perform all our analyses at redshift $z = 0$. The parameters of the assumed cosmologies are given in Table 1 and the simulation parameters in Table 2. (These tables also give parameters for other simulations used in Section 3.)

2.2 Halo definition

We shall be looking at various halo properties defined in spherical shells, so it is appropriate to adopt a spherical overdensity (SO, Lacey & Cole 1994) algorithm to define the haloes. In practice, our halo definition starts by applying a friends-of-friends algorithm (FOF, Davis et al. 1985), with a linking length of 0.2, to construct an initial set of particle groups. We then use the SUBFIND program (Springel et al. 2001) to identify self-bound structures within each group (one of which will always be the main halo itself), as well as the location of the gravitational potential minimum. We identify the potential minimum of the main self-bound structure within each FOF group with the centre of the corresponding dark matter halo. We then grow a spherical boundary around each centre until the total enclosed mass density (not just the original FOF particles) matches that of a virialised halo in the spherical top hat model for a flat cosmology ($\Omega_{\text{tot}} = \Omega_M + \Omega_\Lambda = 1$; see Eke, Cole & Frenk 1996; Bryan & Norman 1998):

$$\Delta_c = \frac{\rho}{\rho_{\text{crit}}} \approx 18\pi^2 + 82(\Omega_M(z) - 1) - 39(\Omega_M(z) - 1)^2. \quad (1)$$

This gives $\Delta_c \approx 94$ for MS and hMS. We shall refer to this halo boundary as its virial radius, R_{vir} , and to the total mass within this radius as M_{vir} .

2.3 Analysis of physical properties

The kinetic and potential energies of each halo (T and U respectively) are computed as in B07, that is,

$$T = \frac{1}{2} \sum_{i=1}^{N_p} m_i v_i^2, \quad (2)$$

where the halo consists of N_p particles; particle i has mass $m_i = m_p$, and velocity vector \mathbf{v}_i relative to the centre-of-mass velocity.

To calculate the potential energy of each halo, we use a random sample of up to $N_{\text{sel}} = 1000$ particles from each halo, and scale the total according to

$$U = \left(\frac{N_p^2 - N_p}{N_{\text{sel}}^2 - N_{\text{sel}}} \right) \left(\frac{-Gm_p^2}{\eta} \right) \sum_{i=1}^{N_{\text{sel}}-1} \sum_{j=i+1}^{N_{\text{sel}}} -W_2(r_{ij}/\eta). \quad (3)$$

We take the form of the potential used in the simulation, which incorporates the SPH smoothing kernel (Springel, Yoshida & White 2001):

$$W_2(x) = \begin{cases} \frac{16}{3}x^2 - \frac{48}{5}x^4 + \frac{32}{5}x^5 - \frac{14}{5}, & 0 \leq x \leq \frac{1}{2}, \\ \frac{1}{15x} + \frac{32}{3}x^2 - 16x^3 + \frac{48}{5}x^4, & \frac{1}{2} \leq x \leq 1, \\ -\frac{32}{15}x^5 - \frac{16}{5}, & x \geq 1. \end{cases} \quad (4)$$

Here, the argument x is given by the ratio of the particle pair separation, r_{ij} , to the spatial softening length, η (see Table 2). We only use the potential (and kinetic) energy of a halo to determine whether or not it is in equilibrium, in order to define the sample of halos to be studied, as explained in the following subsection. The limits we adopt for this criterion are relatively broad so the loss of accuracy introduced by random sampling the halo has no significant impact on the selection. Indeed, when using the smaller, higher resolution simulations in Section 3, we use *all* particles without random sampling, and the results agree well with those in this section.

For computing the specific angular momentum profiles, we divide each halo into a spherical ‘inner’ core region and a series of concentric spherical shells, spaced by 0.2 in $\log_{10}(r/R_{\text{vir}})$, where r is the radial distance from the halo centre. We define an inner region that has radius $0.1R_{\text{vir}}$ for the analysis of angular momentum magnitudes, but radius $10^{-0.6}R_{\text{vir}} \approx 0.25R_{\text{vir}}$ for the analysis of orientations. This larger region is required because the determination of angles between angular momentum vectors is subject to different numerical errors from the determination of the magnitude of the vector, and requires more particles within a given radius to ensure robust and reliable results (see the discussion in Appendix A for full details).

The (cumulative) specific angular momentum vector, $\mathbf{j}(\leq r)$, of the $N_p(\leq r)$ dark matter particles within a given radius r (of total mass $M(\leq r)$) is then given by

$$\mathbf{j}(\leq r) = \frac{1}{M(\leq r)} \sum_{i=1}^{N_p(\leq r)} m_i \mathbf{x}_i \times \mathbf{v}_i, \quad (5)$$

where \mathbf{x}_i and \mathbf{v}_i are the position and velocity vectors of particle i relative to the halo centre and centre-of-mass velocity. The

cumulative specific angular momentum magnitude profile is then given by $j(\leq r) = |\mathbf{j}(\leq r)|$. Since we use the SO algorithm to define our haloes, the total halo specific angular momentum is $\mathbf{j}_{\text{tot}} = \mathbf{j}(\leq R_{\text{vir}})$.

Our choice to investigate the cumulative rather than the differential angular momentum differs from many previous studies (e.g. Bailin & Steinmetz 2005). We made this choice for practical reasons, in order to maximize the size of our sample of well-resolved haloes: if a certain minimum number of particles is required for a reliable measurement of angular momentum, then the use of cumulative profiles implies that this condition only has to be satisfied in the innermost bin considered, rather than in each bin separately. (We discuss our halo selection criteria in detail in the following subsection). Of course, using cumulative profiles has the drawback that the measurements in each bin are not independent, and bin-to-bin variation is reduced compared to the differential case. In practice, however, the angular momentum is dominated by the mass at large radii, so there is very little difference between the behaviour of the cumulative $j(\leq r)$ and the differential $j(r)$. We will show this explicitly in Section 3.

The cumulative angular momentum, $\mathbf{j}(\leq r)$, defined above also differs from the quantity considered in Bullock et al. (2001), the cumulative *mass* profile of specific angular momentum, $M(< j_z)$. Here, j_z is the magnitude of the specific angular momentum projected along the direction of the total angular momentum vector of the halo. Any internal misalignment of the haloes will break the correspondence between this quantity and $j(\leq r)$.

The shapes of the dark matter haloes are computed using the inertia tensor, \mathbf{I} , which directly relates the angular momentum vector, \mathbf{J} , and the angular velocity vector, $\boldsymbol{\omega}$ (i.e. $\mathbf{J} = \mathbf{I} \cdot \boldsymbol{\omega}$). This tensor has components

$$I_{\alpha\beta} = \sum_{i=1}^{N_p} m_i (\mathbf{x}_i^2 \delta_{\alpha\beta} - x_{i,\alpha} x_{i,\beta}) \quad (6)$$

such that $\mathbf{J} = \mathbf{I} \cdot \boldsymbol{\omega}$ (i indexes particles, α and β are the tensor indices with values of 1, 2 or 3, and $\delta_{\alpha\beta}$ is the Kronecker delta). The eigenvectors of the diagonalised inertia tensor define an ellipsoid, which represents the equivalent homogeneous shape of the object in terms of a semimajor axis, \mathbf{a} , intermediate axis, \mathbf{b} and semiminor axis, \mathbf{c} .³

We also calculate the angular velocity magnitude profile, $\omega(r) = |\boldsymbol{\omega}(r)|$. The angular velocity at radius r is defined from the distribution of the mass within that radius through the expression,

$$\boldsymbol{\omega}(r) = \mathbf{I}^{-1}(\leq r) \cdot \mathbf{J}(\leq r). \quad (7)$$

2.4 Halo selection

In their analysis of the halo shapes and spins in the MS, B07 excluded from their halo catalogue haloes that were clearly out of equilibrium at the time of the simulation snapshot. This was achieved by restricting haloes to have an instantaneous ‘virial ratio’ of energies in the range:

$$Q := \left| \frac{2T}{U} + 1 \right| \leq Q_{\text{lim}} \quad (8)$$

³ As noted in B07, this gives the same axes as the mass distribution matrix \mathbf{M} with components $\mathcal{M}_{\alpha\beta} = \sum_{i=1}^{N_p} m_i x_{i,\alpha} x_{i,\beta}$. The two are related through $I_{\alpha\beta} = \text{Tr}(\mathbf{M})\delta_{\alpha\beta} - \mathcal{M}_{\alpha\beta}$ (e.g. Binney & Tremaine 2008).




Figure 1. Mass functions (halo number histograms) for haloes from the four simulations used in this paper. The Millennium Simulation is shown in black with yellow shading, hMS in dashed-black with green shading, DMO in blue and DMG in red. (The DMO and DMG simulations are described in Section 3.) For each simulation, we show the histogram of total mass within R_{vir} (i.e. including stars and gas for the DMG haloes). We show two selections for each simulation: objects that contain at least 300 dark matter particles, within R_{vir} (thin lines, light shading), and within $0.1R_{\text{vir}}$ (heavy lines and shading). In both cases, selected haloes must also satisfy the ‘virialisation’ criterion, $Q \leq 0.5$.

Here, we apply a cut of the same form and adopt the same value of $Q_{\text{lim}} = 0.5$.

B07 found that angular momentum and shape parameters of a halo were subject to numerical biases if it contained fewer than approximately 300 particles. For the analyses presented here, we also adopt the $N_p \geq 300$ selection criterion either for the object as a whole (when profile information is not required), or for the innermost radial bin considered, e.g. $r < 0.1R_{\text{vir}}$. This ensures that only reliable angular momentum profiles are retained and is, by far, the most stringent criterion applied in this paper. The mass functions of haloes selected in this way are shown in Fig. 1.

We apply further selection criteria when investigating the orientation of the angular momentum vectors since the magnitude of the vector is related to the uncertainty in the direction. These are detailed in Appendix A. The number of haloes in each sample are quoted in the legend of the appropriate plots.

2.5 Results

2.5.1 Angular momentum profiles

The cumulative specific angular momentum magnitude profiles, $j(\leq r)$ (as defined in Section 2.3), of haloes in the MS and hMS are shown in Fig. 2. The vertical error bars (directly around the data points) give an estimate of the uncertainty in the median, by analogy to a Gaussian mean,

$$\epsilon_+ = \frac{X_{84} - X_{50}}{\sqrt{N}}, \quad \epsilon_- = \frac{X_{50} - X_{16}}{\sqrt{N}}, \quad (9)$$




Figure 2. Dark matter halo cumulative specific angular momentum profiles, for objects in the MS and hMS simulations. Coloured lines show the median profiles for haloes in different mass bins, and the black lines show the medians for the MS (solid) and hMS (dashed) data together. The error bars on each line (nearly invisible for MS and hMS) are given by Eqn. 9. The outer bars and boxes on the MS line indicate the spread of individual halo profiles around the median (the boxes enclose 68 per cent of the data, the bars enclose 95 per cent). The dotted magenta line shows the $j \propto r$ scaling (with arbitrary normalisation; see text).

where X_i is the value at the i th percentile of the distribution in question, made up of N objects (X_{50} is the median). These are virtually invisible on the lines corresponding to the bulk of the halo populations, but quite significant in the higher mass bin medians which contain considerably fewer haloes. By contrast, the outer bars and boxes (only shown on the MS median line) indicate the spread of the data; the boxes enclose 68 per cent of the data, and the outer bars 95 per cent.

Fig. 2 illustrates the trend of $j(\leq r)$ with mass, as well as the trend with radius at a fixed halo mass. For comparison, we show the radial scaling for a test particle undergoing circular motion at radius r , $j_p(r) = rv_c$, where we take the circular velocity due to the mass within r , v_c , to be constant with radius (as for an isothermal density profile). An NFW density profile (Navarro et al. 1996, 1997) deviates from an isothermal profile at small and large r . Since our haloes are well-described by NFW profiles (Neto et al. 2007; Gao et al. 2008), their angular momentum profiles should deviate slightly from the $j_p \propto r$ form, as indeed they do. Despite the simplicity of these arguments, they give a good account of our simulation results.

A complementary measure of the angular momentum profile is the angular velocity profile, $\omega(r)$. We show this in Fig. 3. We also show the simple scaling for a test particle consistent with the $j_p \propto r$ trend described above: since $v_c = \omega_p(r)r$, we have $\omega_p \propto r^{-1}$. Using ω instead of j removes the mass dependence, but we still find a similar amount of scatter about the median.

Barnes & Efstathiou (1987) showed that the *differential* profile of the specific angular momentum of dark matter halos, $j(r)$, is also roughly proportional to r . Bullock et al. (2001) considered $j_z(r)$, the component of the (differential) specific angular momentum at




Figure 3. Angular velocity profiles, $\omega(r)$ (see Eqn. 7), for haloes in the MS (solid black) and hMS (dashed black) simulations. As with Fig. 2, we plot error bars on both lines, and indicate the spread of the data by the outer bars and boxes on the MS line. The trend from dimensional arguments, $\omega \propto r^{-1}$, is plotted as the magenta dotted line (with arbitrary normalisation; see text).

r parallel to the total halo angular momentum vector; they found that it follows $j_z(r) \propto r^\alpha$, with $\alpha = 1.1 \pm 0.3$. In the (unrealistically) simple case of perfect internal alignment and constant circular velocity throughout the halo, we would expect $j(\leq r)$, $j(r)$, and $j_z(r)$ to all be proportional to r . All these studies see deviations away from this simple trend since haloes are, of course, more complex structures; furthermore, there is a significant amount of variation among individual haloes of a given mass in properties such as mass profiles, spins, and internal alignments.

2.5.2 Spin orientation profiles

We now investigate the orientation of the halo angular momentum vector and its dependence on halo mass and radius. For this, we compute the cumulative angular momentum orientation profile,

$$\cos \theta(\leq r) = \hat{\mathbf{j}}_{\text{inner}} \cdot \hat{\mathbf{j}}(\leq r), \quad (10)$$

where the hat denotes a unit vector (e.g. $\hat{\mathbf{j}} = \mathbf{j}/|\mathbf{j}|$) and the ‘inner’ region is now defined as $r \leq 10^{-0.6} R_{\text{vir}} \simeq 0.25 R_{\text{vir}}$.

The direction of the cumulative angular momentum vector is subject to significant uncertainty due to discreteness effects. To ensure that our results are robust, we have applied a different, additional, set of selection criteria. The net angular momentum of a halo, $\mathbf{j}(\leq r)$, is constructed from the 3-D vector sum of its individual particles’ angular momenta. But haloes, in fact, have very little coherent rotation, so the specific angular momentum of a halo is small compared to the typical specific angular momentum of an individual particle. A halo that has a particularly small j compared to those of its particles will have its direction information dominated by very few particles, which introduces a significant amount of uncertainty. To mitigate this problem, we restrict our sample to haloes whose j values are not too small. We discuss the details of this procedure in Appendix A where we demonstrate that our se-




Figure 4. Cumulative specific angular momentum orientation profiles for haloes from the MS and hMS simulations. The median profile for the entire halo population is shown in black, with error bars and percentile bars shown as in Fig. 2. There is an increasing likelihood of misalignment at larger radius, with a very large amount of scatter between haloes. The coloured lines show the results for the haloes in different mass bins, showing that the lower mass haloes tend to remain better aligned to a larger radius.

lection reduces the scatter and ensures that the angular momentum directions are reliable.

Fig. 4 shows the orientation profiles of the haloes from the MS and hMS simulations. As radius increases, the cumulative angular momentum becomes increasingly poorly aligned with the inner region of the halo. Although the median alignment of the angular momentum at R_{vir} with that in the inner regions is always within 30° , there is a very large scatter amongst haloes. This scatter is much larger than that expected from the numerical limitations discussed above, so we conclude that it reflects the intrinsic variation amongst haloes.

Furthermore, when the data are split into different mass bins, we can see that there is a clear trend, with more massive haloes tending to be less well aligned at large radius. Fig. 5 examines the mass dependence in more detail, by showing how the angle between $\mathbf{j}_{\text{tot}} = \mathbf{j}(\leq R_{\text{vir}})$ and $\mathbf{j}_{\text{inner}} = \mathbf{j}(\leq 0.25 R_{\text{vir}})$ varies as a function of halo mass. While the trend is weak over a wide range of mass at the lower end, there is a clear decrease in alignment for the very highest mass haloes. The reason for this is likely to be related to the hierarchical nature of structure formation: the most massive haloes formed most recently, and are likely to have experienced a major merger more recently than smaller mass haloes. Major mergers have been found to have a strong effect on the halo’s angular momentum magnitude (see e.g. Frenk et al. 1985; Maller et al. 2002; Vitvitska et al. 2002; D’Onghia & Navarro 2007; Zavala et al. 2008); it is highly likely that a major merger would decorrelate the alignment profile of the halo, albeit in a way that could depend strongly on the details of the merger.

B07 considered the alignment of the total halo angular momentum vector with the halo shape axes (their fig. 16). Complementing that, we show the distribution of angles between the inner halo angular momentum and the total halo shape (identified by the




Figure 5. Angle between the specific angular momentum within R_{vir} and within the inner region ($\approx 0.25R_{\text{vir}}$), as a function of halo mass, M_{vir} . There is very little trend for the majority of haloes, but the most massive ones show a clear tendency for increased misalignment. The error bars indicate the uncertainty on the median, and the outer bars and boxes show the spread of the data, as in Fig. 2. The horizontal bars indicate the mass-bin widths.

minor axis, c) in Fig. 6. Although the alignment tends to be good, there is a significant population of haloes where the misalignment exceeds 45° . We find a slight trend with mass which is also likely to arise from the difference in halo merger histories. The distribution of angles between the total halo shape and the inner halo angular momentum is, in fact, very similar to that with the total halo angular momentum shown in B07.

We can attempt to compare our results to those found by Bailin & Steinmetz (2005), hereafter BS05. These authors compared both the angular momentum and shape-axis vectors of haloes as a function of radius. At first glance, our results might seem to suggest better internal halo alignment than found in BS05. However, the differences in halo selection and analysis make direct comparisons complicated. In particular, BS05 used the differential angular momentum profile, $j(r)$, whereas we use the cumulative $j(\leq r)$. Although this allows us to retain more objects, as discussed in section 2.3, it also smooths out quantities like the misalignment profile, resulting in an average reduction in misalignment compared to the differential profile. (This is discussed in further detail in Section 3 where we directly compare differential and cumulative profiles.) While it seems likely that this is the primary reason for the difference between the results, there are other differences in methodology that could also play a role. Although the BS05 simulations have slightly better mass resolution than ours, their simulations had relatively few haloes, which could also increase the scatter compared to our results. The halo populations in the two studies cover a very similar mass range, and we have used similar techniques to control errors due to numerical effects. These include limiting the minimum number of particles in the halo and using bootstrap re-sampling to test for, and reject, haloes whose vector directions are unreliable. In this paper, we also explicitly remove haloes whose energies indicate that they are out of equilibrium, although we do




Figure 6. Normalised histogram of the angle between the halo minor axis and the specific angular momentum of the halo inner region ($\approx 0.25R_{\text{vir}}$). The overall median for the MS and hMS simulations is plotted in black over the coloured lines representing the medians in different mass bins. The error bars show the Poisson uncertainty (\sqrt{N}) in each bin. As well as excluding haloes whose j_{inner} directions are poorly determined due to discreteness, we also apply an analogous selection criteria for the halo shape axes, eliminating the nearly-spherical objects; see Appendix A for details.

not find that this is as strong a constraint as the limits on particle number and those derived from the bootstrap analysis. The lack of a clear trend with halo mass seen in BS05 is consistent with the results we presented in Fig. 5, since the trend we see is quite weak and has a rather large scatter. The similarity in the distributions of angles between the shape axes and angular momentum vectors at R_{vir} and $0.25R_{\text{vir}}$ is also consistent with BS05, given the caveats above. Similar comments can be made regarding earlier work on halo internal alignments, such as that by Dubinski (1992) and Warren et al. (1992), both based on differential profiles. In so far as they can be compared directly, our results are compatible with these.

3 THE EFFECT OF BARYONS

In order to extend our results to realistic galactic haloes, it is necessary to consider how the galaxy formation process affects the dark matter structures. In this section, we investigate how baryonic process modify haloes, in comparison with dark-matter-only structures.

3.1 The simulations

We use a simulation that, at $z = 0$, contains a roughly spherical, high resolution region, with a diameter of about $12.5 h^{-1} \text{Mpc}$, embedded in progressively lower resolution boundary regions out to a cubical boundary of side $35.33 h^{-1} \text{Mpc}$. The simulation was performed with the modified GADGET-2 code developed by Okamoto et al. (2005), which includes a detailed implementation of the physical processes required for forming galaxies.

Although we briefly review these processes here, the reader is referred to Okamoto et al. (2005) for full details. The modelling

of the interstellar medium (ISM) mostly follows the method of Springel & Hernquist (2003): the ISM is a two phase gas, consisting of an ambient hot phase and cool clouds, in pressure equilibrium with each other. The heating and cooling of gas is calculated under the assumptions of collisional ionisation equilibrium in the presence of a uniform ultraviolet background that evolves with time (Haardt & Madau 1996). Cooling depends explicitly on the gas metallicity, using the cooling tables from Sutherland & Dopita (1993); molecular cooling and metal cooling at temperatures below 10^4 K are not included.

Energy to heat the gas is supplied by both Type II and Type Ia supernovae, which provide a feedback mechanism for the gas. Stars form from the cool gas in either a ‘quiescent’ or ‘burst’ mode. In quiescent star formation, stars form according to a given probability once the cold gas density rises above some threshold. These stars have a Salpeter (1955) initial mass function (IMF). In the burst mode, stars form over a shorter timescale and with a top-heavy IMF (Baugh et al. 2005). This gives rise to stronger feedback because of the larger number of supernovae. When the injected heating exceeds the local cooling, the hot gas can be blown out of the galaxy. The starbursts are triggered by the presence of shocks, which are caused by galaxy mergers.

These prescriptions for star formation and feedback ensure that much of the gas remains hot during the early stages of the galaxy formation process and that, once the merger rate has subsided at later times, the gas can cool and form stars stably, resulting in fairly realistic galaxies. We shall refer to the simulation ran with this code as ‘DMG’ (dark matter plus galaxies)⁴.

In addition, we re-ran the DMG simulation without baryons – the ‘DMO’ simulation – redistributing the equivalent baryonic mass to the dark matter particles in the initial conditions to conserve the overall mass. Since these two simulations have exactly the same initial mass distribution, we can match corresponding haloes and thus directly estimate the effects on the halo properties of interest of the baryonic physics included in the simulations. We can then ‘extrapolate’ our conclusions to the MS and hMS, which lack the complex baryonic physics but contain many orders of magnitude more objects.

As in Section 2, we analyse our simulations at $z = 0$. The assumed cosmological parameters are given in Table 1 (they are slightly different to those of the MS and hMS, but this should have no impact on our results), and the simulation parameters are given in Table 2.

We identify haloes in DMO and DMG in the same way as in MS and hMS (see Section 2.2). The spherical overdensity algorithm employs Eqn. 1 to define the ‘virial’ radius of the haloes; for the DMO/DMG cosmology, the overdensity parameter has the value of $\Delta_c = 101$. We have to take account of the low resolution (high mass) boundary particles that surround the high resolution central region in the DMO and DMG simulations. Haloes near the edge of the high resolution region are at risk of contamination by boundary particles. To ensure that this does not affect our results, we retain haloes only if there are no boundary particles within a radius of $R_{\text{vir}} + 100 h^{-1} \text{kpc}$ of their centre.

3.2 Galaxy identification

We identify galaxies as collections of gas and star particles within dark matter haloes. Each galaxy is identified as the most massive object found by a FOF algorithm applied to the baryonic particles (both stars and gas together) within the virial radius of the parent halo, using a linking length of $1.07 h^{-1} \text{kpc}$ (which corresponds to $b = 0.02$ for the stellar particle number density at $z = 0$). This value of b was chosen to be an order of magnitude below the value used for dark matter, since baryons collapse by a factor of $\sim 1/(2\lambda)$, and the spin parameter $\lambda \sim 0.04$ (see e.g. Fall & Efstathiou 1980; Pearce et al. 2001; Baugh 2006). We have checked that this identification is reasonable. The resulting galaxies do not contain either large amounts of hot halo gas or satellite galaxies, which would bias measurements of shape and angular momentum.

The centre of each halo’s galaxy is determined by first finding the centre of mass of the stellar particles within the halo, and the radius of the sphere encompassing them all. The radius is then shrunk by 5 per cent, and the centre of mass of the particles remaining within the sphere is calculated. Using the new centre, the radius is shrunk again, iterating until only ≤ 50 particles remain. The galaxy centre is then taken to be the last centre of mass calculated using at least 50 particles. The galaxy centres are found to correspond very well with the halo centres: for the 99 galaxy–halo systems (with $Q \leq 0.5$ and containing at least 1000 particles in both the halo and stellar component), the median galaxy–halo centre separation, as a fraction of R_{vir} , is 0.0058; the mean is 0.0064 and the standard deviation is 0.0030. We also define the galaxy outer radius to be the distance from the galaxy centre to the farthest baryonic particle included in the FOF group.

3.3 Physical properties

The inclusion of baryons in the DMG haloes requires some definitions to be made more carefully. The virial masses and radii of the DMG haloes are defined using all the mass (i.e. including baryons), as are the kinetic and potential energies. The smaller number of haloes in these simulations makes it possible to compute the potential energies using all the particles within R_{vir} (rather than only using 1000 as in the MS and hMS simulations). We have to take into account the different gravitational softening lengths and individual masses of the baryonic particles, so the contribution to the potential energy from each particle pair i – j is the mean:

$$u_{ij} = \frac{1}{2} \left(\frac{-W_2(r_{ij}/\eta_i)}{\eta_i} + \frac{-W_2(r_{ij}/\eta_j)}{\eta_j} \right), \quad (11)$$

where $W_2(x)$ is the SPH smoothing kernel (see Eqns 3 and 4). The total potential energy of each halo system is therefore:

$$U = -G \sum_{i=1}^{N-1} \sum_{j=i+1}^N m_i m_j u_{ij}. \quad (12)$$

We perform a dynamical decomposition of the stellar systems in our simulations in order to determine the relative contributions of the disc and bulge components, measured both according to stellar mass and to B -band luminosity. We do this for galaxies containing at least 5000 star particles, in the same way as was done by Okamoto et al. (2005), whose method is based on that of Abadi et al. (2003b). The stringent particle number criterion leaves only 30 galaxies in the sample but it ensures that we exclude galaxies whose morphologies are biased towards bulginess due to poor numerical resolution. We let the angular momentum of the stellar

⁴ The DMG simulation was also analysed in Libeskind et al. (2007), (their ‘SR’ simulation) to investigate satellite galaxy alignments.




Figure 7. Disc-to-total ratios (D/T) for galaxies containing at least 5000 star particles (absolute B -band magnitude brighter than -17.9), in haloes containing at least 1000 particles, as a function of the number of star particles in each galaxy. The upper panel shows the number of particles as a function of D/T and the lower panel shows the distribution of this quantity. We plot results for the D/T ratio according to stellar mass (red histogram and crosses) and according to B -band luminosity (blue histogram and asterisks). The constraint on the number of star particles reduces the sample to 30 galaxies, but their D/T values do not depend significantly on the number of particles. Imposing the virial ratio limit $Q \leq 0.5$ has no effect on the selected sample.

component of each galaxy define a ‘ z ’-axis, and compute the component of the angular momentum of each star particle parallel to this direction. Half of the bulge is identified with the particles that have $j_z < 0$; the total bulge mass is defined as twice the mass of those particles. The disc mass is then given by the difference between the total stellar mass and the bulge mass. Using the same method, we also compute the disc-to-total ratio (D/T) in terms of the star particles’ B -band luminosity (see Okamoto et al. (2005) for details).

The distribution of D/T ratios for galaxies in our simulations is shown in Fig. 7. The faintest galaxy in the sample has a total B -band luminosity corresponding to an absolute magnitude of -17.9 . This distribution does not depend significantly on the number of star particles in the galaxy. Two thirds of our sample have B -band D/T ratio greater than 0.5. Thus, our simulation produces a distribution of morphological types that is broadly consistent with observations (Benson et al. 2007; Parry et al. 2009). This is a significant success of our simulations.




Figure 8. Distribution of the mass ratios of matched haloes in the DMO and DMG simulations. The haloes in each simulation are selected to have $Q \leq 0.5$, and at least 300 dark matter particles within $0.1R_{\text{vir}}$. Matching the haloes in the two simulations involves finding each halo’s closest counterpart within $100 h^{-1}\text{kpc}$. The vast majority of matched halo pairs have very similar masses to each other.

3.4 Results

The basic halo selection for the DMO and DMG haloes is the same as for MS and hMS described in Section 2.4: we require haloes to have $Q \leq 0.5$, and $N_p(\leq r) \geq 300$ within the innermost radius considered. Using $0.1R_{\text{vir}}$ as the innermost radius yields 83 haloes from DMO and 88 haloes from DMG. Again, we use slightly different selection criteria for our investigation of angular momentum orientations (see Appendix A).

In order to compare the same haloes in the simulations with and without baryons, we match each selected dark matter halo in DMG with its counterpart in the set of selected DMO haloes, by finding the closest DMO halo centre within $100 h^{-1}\text{kpc}$ of the DMG halo centre. Out of the 83 and 88 haloes selected as described above, clear matches are found for 67 halo pairs. Fig. 8 shows the distribution of their mass ratios. The few unmatched haloes do not have a counterpart within $100 h^{-1}\text{kpc}$ that passes the selection criteria. This failure to match all of the haloes is a feature of the simulations, not of the matching algorithm – it does not depend on the selection criteria or the limiting separation. Indeed, the fact that the same selection criteria yield differing halo numbers for the DMO and DMG simulations hints at this. This difference is likely to be due to the baryons having a significant influence on the evolution of some haloes, causing some to merge in one simulation but not in the other, or occasionally causing some haloes to have very different numbers of dark matter particles.

3.4.1 Angular momentum profiles

We show the cumulative specific angular momentum profiles of haloes in the DMO and DMG simulations in Fig. 9. These show the same basic trends that we saw in the MS and hMS simulations (Figs. 2 and 3; we also plot the same $j \propto r$ line for reference). The mass dependence of $j(\leq r)$, present in the top panels, is re-




Figure 9. Cumulative radial profiles of: specific angular momentum (top), scaled angular momentum $j(\leq r)/(V_{\text{vir}} R_{\text{vir}}) = j(\leq r)/\sqrt{GM_{\text{vir}} R_{\text{vir}}}$ (middle) and angular velocity (bottom), for the dark matter in haloes from the DMO and DMG simulations (left and right panels respectively). The profile for each halo is colour coded by its total mass (M_{vir}). The heavy black lines are the median profiles, with error bars calculated according to Eqn. 9. As a guide, the dotted magenta lines show the radial scaling derived from a simple argument assuming constant circular velocity, with arbitrary normalisation (see text). Scaling the angular momentum of each halo by its circular velocity and radius removes the mass dependence of $j(\leq r)$.

moved by scaling by $V_{\text{vir}} R_{\text{vir}} \equiv v_c(R_{\text{vir}}) R_{\text{vir}}$, where the circular velocity $v_c(r) = \sqrt{GM(\leq r)r^{-1}}$. We plot $j(\leq r)/(V_{\text{vir}} R_{\text{vir}})$ in the middle panels of Fig. 9. The mass dependence is also absent when plotting the angular velocity, $\omega(r)$. There is a similar degree of halo-to-halo scatter as in the MS and hMS simulations.

For completeness, and to aid comparisons, we have also computed the *differential* specific angular momentum profiles for the DMO and DMG haloes. The cumulative and differential specific angular momenta are related through:

$$j(\leq r) = \frac{\int_0^r j(r') \rho(r') r'^2 dr'}{\int_0^r \rho(r') r'^2 dr'}. \quad (13)$$

If $j(r) \propto r$, and the density profile $\rho(r) \propto r^{-2}$ (at least, for the radius and mass ranges of interest), then the cumulative specific angular momentum is simply $j(\leq r) = j(r)/2$. In Fig. 10, we plot the medians of both $j(\leq r)$ and $j(r)/2$, scaled by $V_{\text{vir}} R_{\text{vir}}$ (as in the middle panels of Fig. 9) to remove the dependency on halo mass. The cumulative and differential angular momentum profiles are very similar, but there are far fewer halos included in the differential calculation.

Our selection criteria for accurate angular momentum measurements requires that, in the differential case, *each* bin should




Figure 10. Comparison of the cumulative and differential specific angular momentum profiles. In the top panels, we plot the median of the cumulative $\log_{10} j(\leq r)/(V_{\text{vir}} R_{\text{vir}})$ (blue; same as the middle panels of Fig. 9), and the differential $\log_{10} j(r)/(2V_{\text{vir}} R_{\text{vir}})$ (red; see text). The differential data are plotted at the centres of their radial bins, with the cumulative data plotted at the upper ends. For the differential quantities, the bin that ends at $0.1 R_{\text{vir}}$ has the same width in $\log_{10} r$ as the other bins, rather than encompassing all mass down to the centre. To illustrate the scatter in the data, we plot boxes and outer bars enclosing the 68 and 95 per cent of the distributions in each bin respectively. The lower panels show the fraction of haloes in each radial bin that have at least 300 dark matter particles, relative to the number selected for the cumulative calculation.

satisfy the conditions described in Section 2.4. In particular, each bin has to contain at least 300 dark matter particles. In the lower panels of Fig. 10, we show the fraction of the 83 and 88 halos selected for the cumulative calculation that are also selected for the differential calculation, i.e. that have at least 300 dark matter particles in each radial bin. Because of the relative lack of material in the inner bins, there is a significant reduction in the number of selected haloes in the differential case (of course, preferentially excluding those with lower masses).

Fig. 11 shows the median value of $\omega(r)$, the angular velocity at radius r due to the mass internal to that radius, in each of our four simulations. It is immediately apparent that the dark matter in the inner regions of haloes containing galaxies spins significantly faster than in their dark matter-only counterparts. This can be seen even more clearly in Fig. 12, where we plot the ratio of $j(\leq r)$ for each DMG halo to that of its DMO counterpart. Note that we have fewer halos after matching, as discussed above. (The equivalent plot for ratios of $\omega(r)$ is very similar.) There is a large halo-to-halo scatter but, at the virial radius, the dark matter remains unaffected by the baryons (in the median). However, within R_{vir} , the trend in the median is for the angular momentum of the inner 10 per cent of the dark matter in a halo to be ~ 50 per cent greater if that halo has a galaxy in it than if it does not.

There are two possible reasons for the difference between the j_{inner} of the DMG and DMO haloes. The addition of baryons can cause the dark matter to contract adiabatically, so that a given mass is contained within a smaller radius, and therefore rotates faster. In this case, the angular momentum of the dark matter would be




Figure 11. The median dark matter angular velocity profiles for haloes in our four simulations. The profiles for haloes containing baryons (DMG) have a significantly greater angular velocity in their inner regions, in the median. As in previous figures, the error bars give the uncertainty in the median (Eqn. 9) and the outer bars and boxes plotted on the MS curve give the spread of the data (the boxes enclose 68 per cent, and the outer bars 95 per cent of the data).

conserved independently of that of the baryons. Alternatively, the dark matter could also gain angular momentum through the transfer from the infalling gas by tidal torques and dynamical friction. We now look into these possibilities in more detail.

We expect a DMG halo to be more centrally concentrated than its DMO counterpart because its baryons will have dissipated their energy, fallen to the centre and deepened the potential well (see e.g. Barnes & White 1984, Blumenthal et al. 1986, Gnedin et al. 2004). Dark matter will fall into the deeper potential well and, if it conserves its own angular momentum, it would spin up, giving rise to an increase in $j(\leq r)$ in the central regions. Thus, if the angular momentum of the dark matter of a DMG halo within a given inner radius is the same as that of the dark matter of its DMO counterpart within some *larger* radius, which contains the same dark matter mass, then there has been no net transfer of angular momentum between the dark matter and the baryons: the dark matter has simply contracted (perhaps adiabatically).

Let M'_{DMG} be the mass of dark matter contained within $0.1R_{\text{vir}}$ of the DMG halo, appropriately scaled by the baryon fraction⁵, and let r_0 be the radius in the corresponding DMO halo that contains the same mass, i.e:

$$M'_{\text{DMG}}(\leq 0.1R_{\text{vir}}) = M_{\text{DMO}}(\leq r_0) \quad (14)$$

(If the dark matter has contracted substantially, we would expect $M_{\text{DMO}}(\leq r_0) > M_{\text{DMO}}(\leq 0.1R_{\text{vir}})$.) We then compute the ratio between the dark matter angular momentum, $j_{\text{DMG}}(\leq 0.1R_{\text{vir}})$, and $j_{\text{DMO}}(\leq r_0)$; if the dark matter angular momentum has been




Figure 12. Logarithm of the ratio of the cumulative specific dark matter angular momentum profiles of haloes from the DMG simulation to that of the matched haloes in the DMO simulation. Each line represents the ratio from a matched DMG-DMO halo pair, colour coded according to the mean mass of the two haloes. The median ratios are marked by the heavy black line, with error bars given according to Eqn. 9.

conserved, this should be unity. In Fig. 13, we compare this to the ratio of dark matter masses within $0.1R_{\text{vir}}$. (We find that, in practice, $R_{\text{vir,DMG}} \simeq R_{\text{vir,DMO}}$, so we do not distinguish them in this plot.) Fig. 13 shows that the increase in angular momentum of the inner dark matter in DMG haloes at fixed mass (median of 60 per cent) is much greater than the corresponding change in mass at that radius (median value of 10 per cent). The fact that j grows even at fixed dark matter mass suggests that the angular momentum of the dark matter is not conserved, but instead gains, at least in part, from the baryons. This is consistent with the results of Kaufmann et al. (2007), who performed a detailed investigation of the different ways in which angular momentum can be transported away from the cooling gas in a halo, given that current simulations have not yet reached the resolution whereby the final angular momentum of the gas has converged.

Our work builds upon the results of Zavala, Okamoto & Frenk (2008), which was also based on simulations by Okamoto et al. (2005). They found that in a simulation with weak star formation and feedback, a large amount of angular momentum was transferred from the baryons to the dark matter, resulting in a bulge-dominated galaxy. On the other hand, in a simulation from *the same initial conditions* but stronger feedback, the baryons approximately conserved their angular momentum resulting in a disc-dominated galaxy. The simulations that we analyse here use the same strong feedback model. Our results for a sample of many objects indicate that, in fact, some angular momentum transfer still takes place on average even in the strong feedback case. The size of this transfer, however, is sufficiently small and has sufficiently large scatter that the majority of the galaxies still end up being disc-dominated.

⁵ When comparing the dark matter mass between the DMO and DMG haloes, we scale each DMG halo dark matter mass by f_{bary} , the fraction of mass in baryons within R_{vir} ; i.e. we set $M'(\leq r) = M(\leq r)/(1 - f_{\text{bary}})$.




Figure 13. Ratio of the specific angular momentum of the dark matter mass within $0.1R_{\text{vir}}$ of each DMG halo, to that of the corresponding DMO halo at a radius r_0 that contains the same mass. A value of unity (dashed horizontal line) indicates that the dark matter has conserved its own angular momentum. This ratio is plotted against the ratio of dark matter masses at $0.1R_{\text{vir}}$ for each halo pair. (The DMG dark matter mass is scaled by the baryon fraction, see text). Values greater than unity (dashed vertical line) in this axis indicate that the halo has contracted. To highlight the relative scale, the 1:1 relation is marked with a dotted line. Haloes are selected according to our usual criteria of $N_p(\leq 0.1R_{\text{vir}}) \geq 300$ and $Q \leq 0.5$. Blue circles mark the 29 halo pairs with $r_0 \geq 0.1R_{\text{vir}}$; red crosses mark the 34 haloes where $r_0 < 0.1R_{\text{vir}}$. In the latter case, the measured $j(\leq r_0)$ could come from fewer than 300 particles, but in practice this only occurs for one halo (marked with a magenta circle around its red cross) since most have $r_0 \approx 0.1R_{\text{vir}}$. The median of all the halo pairs excluding this one is marked with a black cross, with error bars given by Eqn. 9.

3.4.2 Spin orientation profiles

We examine the dark matter cumulative spin orientation profiles of haloes in the DMO and DMG simulations in the same way as we did earlier for the MS and hMS data; note that the ‘inner’ region is again set at $\approx 0.25R_{\text{vir}}$. From the individual halo profiles shown in Fig. 14, we can see that, as with the MS and hMS haloes, there is a trend in the median such that the dark matter angular momentum vector at R_{vir} is $15\text{--}30^\circ$ away from that of the inner dark matter, and also that there is a very large scatter about that trend. There is no discernible trend with halo mass; given the small mass range of these haloes, however, this is consistent with the results shown earlier (Fig. 5). When comparing the median trends of DMO to those of DMG (lower panel of Fig. 14), there is a suggestion that the haloes that have experienced baryonic physics have a total spin that is slightly more closely aligned than the dark-matter-only haloes. However, the two curves are within each others’ error bars, so this result on its own is inconclusive.

We have also computed the *differential* orientation profiles of the dark matter angular momentum. As for the differential angular momentum magnitude profiles (Fig. 10), the results are qualitatively similar to the cumulative case, but there is now much more scatter. This is due both to the fact that the points in each bin are now independent, and to the reduction in the number of haloes selected. The restriction that each bin must contain at least 1000 dark




Figure 14. Cumulative dark matter angular momentum orientation profiles of haloes in the DMO and DMG simulations (top left and right respectively). Each halo is colour coded according to its mass, with the median profile shown in heavy black, and error bars plotted according to Eqn. 9. The bottom-left panel re-plots just the two median lines and their error bars, with DMO and DMG shown in blue and red respectively.

matter particles leaves us with just 21 and 18 haloes from the DMO and DMG simulations respectively. The median misalignment between the inner dark matter and that at R_{vir} is increased to $30\text{--}35^\circ$.

In order to examine possible differences between their alignment distributions in more detail, we compare the cumulative orientation profiles of DMG haloes directly with their counterparts in the DMO simulation. These results are shown in Fig. 15. There is a definite tendency for the baryonic processes inside the DMG haloes to change the orientation of the inner dark matter angular momentum, while the outer regions of the haloes remain well aligned with their DMO counterparts. We conclude that the baryons tend to cause the inner regions of haloes to become better aligned with their total halo angular momentum vector. This conclusion is reinforced by the tentative results from Fig. 14.

This result is in agreement with that of Bailin et al. (2005) who calculated the angle between the minor axis of haloes simulated with and without galaxy formation physics. They found that the inclusion of baryonic processes reorients the inner halo shape axes while leaving the outer halo unchanged.

3.4.3 Dark matter–galaxy alignment distributions

For most practical purposes, the orientation of a halo relative to its galaxy is more relevant than the intrahalo orientation profile itself. The alignment of a galaxy with its halo is, in principle, an observable quantity, since gravitational lensing of background objects can be used to measure the size and shape of the mass distribution surrounding a galaxy (although in practice this is extremely difficult). The galaxy–halo orientation is also an important input to semianalytic models of galaxy formation which, with this information in hand, could be used, for example, to construct mock galaxy catalogues relevant to lensing studies.

As shown in previous work (on these same simulations, Libe-




Figure 15. Angle between the cumulative specific angular momentum vectors of dark matter in the DMG simulation and their counterparts in the DMO simulation. (A value of $\cos \theta = 1$ means that the dark matter mass within that radius in the DMG halo is aligned with that in the DMO simulation.) Each halo pair is colour coded according to the mean of the two halo masses, and the median trend with error bars is marked with the heavy black line.

skind et al. 2007), the galaxy angular momentum is a very accurate proxy for the orientation of the galaxy itself (i.e. the minor axis of its mass distribution)⁶. Therefore, in this section, we shall use the angular momentum of the stellar component of the galaxies (j_{gal}), in addition to the minor axis of the stellar mass distribution (c_{gal}), to define their orientation.

Fig. 16 shows the median radial profile of the orientation of the galaxies in the DMG simulation with respect to the angular momentum vector of the dark matter of either their parent halo in DMG, or of the corresponding halo in DMO. The scatter in the angles is very large and so we do not show the usual percentile boxes and bars; instead the error bars give the uncertainty in the median. Only a rather weak trend with radius is apparent: the median values increase from 25° at $\approx 0.25R_{\text{vir}}$ to 35° at the virial radius, but given the large scatter, this is of very low significance.

We examine the angular momentum alignment distributions themselves in Figs. 17 and 18. These show histograms of the cosine of the angle between j_{gal} and the dark matter j_{tot} and j_{inner} respectively. There is a clear concentration towards small angles, particularly for the inner halo, but both distributions have a long tail towards large angles. Furthermore, the distributions are very similar for the DMG and DMO haloes. This relatively weak galaxy–halo alignment serves to wash out the subtle changes in the relative orientations of the dark matter from the DMO and DMG simulations seen in the previous section.

Fig. 19 shows the histogram of the angle between the full halo




Figure 16. Orientation profiles of the galaxies in the DMG haloes with respect to the cumulative dark matter angular momentum vector of the DMG parent haloes (red), and the corresponding DMO haloes (blue). Only the median profiles are plotted. Note that the error bars are the uncertainties on the median (Eqn. 9), *not* the spread of the data, which is much larger.




Figure 17. Distribution of the angle between the specific angular momentum vector of the galaxies in DMG haloes and the specific angular momentum of the total dark matter of the DMG halo itself (red) or the corresponding DMO halo (blue). The medians of the distribution are marked (at arbitrary heights) with error bars given by Eqn. 9.

⁶ We find the median cosine of the angle between the minor axis and angular momentum axis of the stellar components of a galaxy is $0.9949^{+0.00047}_{-0.00443}$. (This uses the 81 galaxies with ≥ 1000 stellar particles, and stellar shape axis ratios of $s = c/a \leq 0.81$, in haloes with $Q \leq 0.5$ and at least 1000 dark matter particles.)

minor axis, c_{tot} , and the galactic spin axis. In this case, the scatter between the orientation of the halo angular momentum and its mass distribution (Fig. 6 and B07) tends to wash out most of the (already weak) alignment between galaxy and halo. Finally, we consider the galaxy–halo misalignment purely in terms of the shape of the system: we plot the distribution of the angle between the halo minor axis and the galactic stellar minor axis in Fig. 20. (Although we




Figure 18. As Fig. 17, but this time for the angle between the *inner* halo dark matter angular momentum vector and the central galaxy. The medians and their uncertainties are marked.




Figure 19. As Fig. 17, but for the angle between the halo minor axis c_{tot} and the spin axis of the central galaxy. Since there is no distinction between parallel and antiparallel for the halo shape axis, we plot the absolute value of the dot product.

no longer need to select by the angular momentum magnitude, we do have to exclude galaxies that are overly spherical, just as we do for the haloes, which reduces our sample size further.) The alignment distribution is still very broad, but there is a slightly more pronounced peak towards alignment. A large number of galaxies are oriented perpendicular to the halo and half of them are misaligned by 45° or more. In Fig. 21, we show the distribution of the angle between the galaxy and the *major* axis of the halo (a_{tot}). About 10 per cent of the galaxies have their minor axes aligned within 30° of




Figure 20. As Fig. 19, but for the angle between the minor axis of the halo c_{tot} , and that of the central galaxy c_{gal} . There is no selection by j , but galaxies are selected by their axis ratio, taking the same limit as the halo shape.




Figure 21. As Fig. 20, but for the angle between the halo major axis a_{tot} and the minor axis of the central galaxy.

their halo's major axis. In the next subsection, we will explore the consequences of these results for gravitational lensing studies.

We have also examined the alignment distributions when splitting the galaxy populations according to their disc-to-total ratios (see Section 3.3). The resulting galaxy samples are small, however, and we find no significant dependence of galaxy–halo alignment on galaxy morphology.

In their analogous work on halo shape axes, Bailin et al. (2005) found that the presence of baryons forced the inner halo to align itself with the galactic disc. We do not find such a strong

correlation. This may be due to a combination of factors including the different physical models adopted in the two studies and, importantly, the larger sample size in our simulations. Our results are consistent with those obtained for a sample of four haloes by Gustafsson et al. (2006), both as regards the variability from object to object and the overall trend.

Croft et al. (2009) measured the alignments between shape and angular momentum for the different components of their galaxies (at $z = 1$), defined as the dark matter and the gas and stars in the self-bound subhaloes in their simulation. Their results show slightly worse alignment than we find (median angle of 43.5° between dark matter and stars, compared with our values of 23.9° and 34.4° for the median angle between the galaxy and the inner and total halo respectively). The difference in the type of objects considered in each study is significant; we only consider the central galaxy in each halo, as our simulations do not have enough resolution to study subhaloes properly. We also take great care to remove objects whose combination of angular momentum size and particle count would lead to large errors in the determination of orientation. Including such objects would have introduced greater scatter in the results. Croft et al. (2009) do perform resolution tests, however, which are very informative. In their lower resolution run, the median alignment is significantly poorer. The resolution in our simulations is slightly better than theirs, so our results could be consistent.

In recent years, studies based on the SDSS and 2dFGRS have shown that satellite galaxies in haloes are preferentially aligned along the major axis of their central galaxy (Brainerd 2005; Yang et al. 2006; Azzaro et al. 2007; Wang et al. 2008; Sales & Lambas 2004, 2009). By comparing with simulations, Agustsson & Brainerd (2006) and Kang et al. (2007) showed that if the central galaxies were oriented such that their angular momentum axis were aligned with their halo’s minor axis or angular momentum axis, then the resulting satellite distribution would be preferentially aligned with the central galaxy’s major axis. In fact, the latter paper showed that if the galaxy and halo minor axes were perfectly aligned, then the resulting satellite alignment signal would be much stronger than observed. A misalignment by $\sim 40^\circ$ on average, such as is produced by orienting the galaxy minor axis with the halo angular momentum, is required to bring the signal down to match the observations. Wang et al. (2008) have recently shown that this result applies to haloes with masses almost down to those in our DMO/DMG simulations. Fig. 20 shows that our galaxy–halo (mis)alignments are fully consistent with the picture presented in these studies, with a median angle between galaxy and halo minor axis of $\sim 45^\circ$. Libeskind et al. (2007) used this DMG simulation to study the alignment of satellite galaxies, and found results consistent with the observations, although there were only three sufficiently well-resolved systems of satellites to study.

3.4.4 Alignment of the projected mass distribution

The misalignment between haloes and their galaxies can have very important consequences for observational attempts to measure halo properties, particularly from gravitational lensing data. In practice, the lensing signal from an individual galaxy halo is too weak to be useful; the lensing distortion has to be averaged over many galaxies by stacking data from appropriately scaled and aligned images (Natarajan & Refregier 2000). Thus, in order to make predictions for the observable shapes of dark matter haloes from simulations, we need to consider stacked projected shapes, rather than the full 3-D triaxiality/sphericity distribution (as studied in e.g. B07).

The broad distribution of galaxy–halo (mis-)alignments (Figs. 17, 18, 19, 20 and 21) will have a significant impact on estimates of halo shape from stacked 2-D mass distributions. To quantify this, we first consider the population of relaxed, well-resolved DMG haloes containing well-resolved galaxies (i.e. $Q_{\text{lim}} = 0.5$ and at least 300 dark matter particles in each halo, as well as at least 300 star particles in each galaxy). Since we have access to the full particle coordinates, we can rotate each halo in 3-D such that the central galaxies’ major–minor axes planes are all aligned. We then compute the 2-D projected mass distribution matrix, \mathbf{M} , of the halo, which has components:

$$\mathcal{M}_{\alpha\beta} = \sum_{i=1}^{N_p} m_i r_{i,\alpha} r_{i,\beta}, \quad (15)$$

where the sum is over all particles in the halo (dark matter and baryons), and α, β denote the matrix indices (1 or 2), such that $r_{i,1}$ is the halocentric distance of particle i in the direction parallel to the galaxy’s major axis and $r_{i,2}$ is the distance parallel to the galaxy’s minor axis. These matrices are then normalised by the halo size, $\mathbf{M}' = \mathbf{M} / (M_{\text{vir}} R_{\text{vir}}^2)$.

Of course, this luxury – rotating in 3-D, then projecting – is not available to observers, who are instead restricted to aligning the projected galaxy images in 2-D. Our method is not wholly realistic therefore, but it does give an idealised ‘best-case’ setup, by allowing us to ‘observe’ each galaxy–halo system with the galaxy edge-on. Furthermore, in practice, observers are likely preferentially to select galaxies which appear more edge-on (e.g. Hoekstra et al. 2004; Parker et al. 2007).

For each halo, we compute the eigenvectors (\mathbf{a} , \mathbf{b}) and eigenvalues ($a > b$) of \mathbf{M}' , giving the distribution of projected halo shapes. Finally, we take the mean of \mathbf{M}' over all the selected galaxies, and compute its eigenvectors and eigenvalues. These then describe the net shape distribution of the selected haloes, in a 2-D projection, (perfectly) scaled and aligned by their galaxy.

Fig. 22 shows the ellipse defined by the eigenvectors of the stacked system, along with those from each individual halo. Here we can see that although the projected mass distribution of an individual halo is by no means necessarily circular, the misalignment of the haloes with their galaxies results in a stacked mass distribution that is almost exactly circular: the axis ratio is $s = b/a = 0.989$ (so the eccentricity is $\epsilon = \sqrt{1 - b^2/a^2} \simeq 0.1$ and the ellipticity is $e = 1 - b/a \simeq 0.01$). Given the uncertainties associated with measuring the directions of the shape axes, this is indistinguishable from a circular distribution.

We can also label the projected halo axes according to which galaxy axis they are closest to. The eigenvector lying closer to parallel to the galactic major axis we label α , and that closer to the galactic minor axis we label β . Thus, we have

$$\alpha = \begin{pmatrix} \max(a_x, b_x) \\ \min(a_y, b_y) \end{pmatrix}, \quad \beta = \begin{pmatrix} \min(a_x, b_x) \\ \max(a_y, b_y) \end{pmatrix}, \quad (16)$$

where the x and y subscripts refer to the galactic major and minor axes respectively (see Fig. 23 for a diagram). Orthogonality means that the axis ratio $\zeta := \beta/\alpha$ is equal to either s or $1/s$, depending on the galaxy–halo orientation. The distributions of the individual and stacked axis ratios themselves (both s and ζ) are shown in Fig. 24. Of the 171 haloes in the distribution, 77 (45 per cent) have $\zeta > 1$, i.e. have their projected halo major axis more aligned with the minor axis of their galaxy. Even though the distributions peak significantly away from $b/a = 1$, the stacked result (marked by the arrow) is essentially indistinguishable from unity. We have esti-




Figure 22. The 2-D projected normalised mass distributions of DMG haloes (coloured dotted lines). Each halo has been aligned according to the orientation of its galaxy’s stellar mass distribution. The colouring is by the total halo mass, as indicated in the bar at the top. The net result from stacking all haloes is the heavy black ellipse, with arrows marking the semimajor (heavy) and semiminor (lighter) axes. Plotted beneath the stacked mass distribution ellipse is a circle (heavy magenta) as a visual aid. The axes are labelled in the dimensionless units of the matrix \mathbf{M}' (see text).




Figure 23. Diagram illustrating the definition of the α and β projected halo shape axes, compared to the standard major and minor axes (a and b ; see Eqn. 16). The ellipse of a projected halo mass distribution is plotted onto axes that are aligned with its central galaxy. We show two different possible orientations of the halo with respect to the galaxy: on the left, the halo and galaxy axes are nearly aligned, and on the right they are nearly perpendicular. These two orientations show the two possible ways in which α and β are assigned to a and b .

mated the error on this stacked halo result, by bootstrap resampling the projected halo shape data, and recomputing the stacked result. The two percentile boxes and bars in Fig. 24 show the median and spread (68 per cent of the data within the boxes, 95 per cent within the bars) of the data from 5000 bootstrap resamplings.

There have been several attempts to measure the shapes of dark matter haloes from weak lensing data. Hoekstra et al. (2004) claimed to have found a significant detection of halo ellipticity, excluding circular haloes at the 99.5 per cent confidence level, and




Figure 24. Normalised histogram of the axis ratios s (red) and ζ (blue) of the 2-D projected mass distributions of haloes shown in Fig. 22. Unity is marked with a dashed line. The axis ratio of the stacked mass distribution is indicated by the arrow near $b/a \simeq 1$. Below the arrow are two points (s in red, ζ in blue, arbitrary heights) with percentile boxes and bars (as in Fig. 2), showing the spread of data from bootstrap resampling the projected halo distribution (see text). This gives an estimate of the uncertainty on the stacked halo result.

yielding an average axis ratio of ≈ 0.7 . However, they lacked redshift data and this limited the accuracy of their method. Parker et al. (2007) obtained a similar result (an axis ratio of ~ 0.7 at a 2σ level) from the CFHT Legacy Survey, again without redshift data. Mandelbaum et al. (2006) used the large dataset of the SDSS (including redshifts and morphologies), and performed a very thorough exploration of possible systematic effects (see also Mandelbaum et al. 2005). Their results were far less conclusive than the previous studies, with no definitive detection of halo ellipticity. There was a hint that spiral galaxies are aligned perpendicular to their haloes when averaged over luminosities, and that ellipticals are increasingly well aligned with increasing luminosity. However, these results were not statistically significant. We have checked if our results depend on galaxy morphology, as defined by the disc-to-total ratio discussed earlier: We are not able to detect any statistically significant variation, although our sample sizes in this case are small.

These studies highlight the practical difficulties in inferring halo properties from weak lensing data. The results we have presented here show that the expected signal of non-circularity is extremely weak. When assessing this conclusion, two caveats must be born in mind. Firstly, the SO algorithm with which we identify haloes inherently biases the halo mass distributions towards sphericity (see B07). However, Figs. 22 and 24 show that this is not the main cause of the near spherical shape of the stacked halo. Rather, it is the misalignment between galaxy and halo that smears out the significant projected ellipticities of the haloes.

Secondly, we have only shown the projected halo ellipticities at the virial radius. Other studies have shown that CDM haloes become increasingly aspherical towards the centre (Hayashi et al. 2007). However, it is not clear that this is still the case for systems with baryons: the limited studies so far of the shape profiles of CDM haloes with baryons suggest that, regardless of mass,

haloes could become *more* spherical towards the centre, depending on their merger histories (Kazantzidis et al. 2004; Bailin et al. 2005; Gustafsson et al. 2006; Berentzen & Shlosman 2006; Abadi et al. 2009).

4 CONCLUSIONS

In this paper we have investigated the angular momentum structure of dark matter haloes in the Λ CDM cosmology. Our first and largest sample includes $\sim 10^5$ well resolved haloes of galaxy, group and cluster masses drawn from the Millennium Simulation and from a smaller, higher resolution dark matter simulation (hMS). Our second sample of ~ 50 –100, also well resolved, galactic-size haloes is drawn from two smaller simulations ran from the same initial conditions, one with only dark matter and the other with baryons as well (which can undergo cooling, star formation, feedback, etc.)

In the simulations without baryons, we have found that the median of the cumulative specific angular momentum of dark matter as a function of radius, $j(\leq r)$, scales approximately as $j \propto r$ (i.e. the scaling that corresponds to circular motion in a mass distribution with constant circular velocity). The amplitude of this trend scales with halo mass, but even at a fixed mass, the scatter is large (over an order of magnitude). The dependence on halo mass is removed if $j(\leq r)$ is scaled by $V_{\text{vir}} R_{\text{vir}}$, i.e. $j(\leq r)/\sqrt{GM_{\text{vir}} R_{\text{vir}}}$. The angular velocity, $\omega(r)$, is similarly independent of mass. These quantities, however, still exhibit large scatter. Thus, haloes do not rotate like solid bodies, but have a rotation structure close to $\omega \propto r^{-1}$. These results apply over the ~ 5 orders of magnitude in halo mass spanned by our pure dark matter simulations.

We investigated the coherence of the halo angular momentum by studying how the orientation of the cumulative angular momentum vectors at different radii deviate from that at an ‘inner’ radius of $\approx 0.25 R_{\text{vir}}$, j_{inner} . In the median, the total specific angular momentum (i.e. of the mass within virial radius) is directed about 25° away from j_{inner} . Again there is large scatter: Ninety-five per cent of the haloes have their total angular momenta directed between 5° and 65° away from their j_{inner} . There is a weak trend with mass, with the more massive haloes showing a greater range of alignment angles (and greater misalignment in the median). In the hierarchical model of structure formation, the more massive haloes are more likely to have experienced a recent merger event that could have altered their angular momentum structure.

We then investigated the effects that baryons have on the rotational structure of haloes in the simulations described by Libeskind et al. (2007), using the techniques introduced by Okamoto et al. (2005). Although the volume of this simulation (diameter $\sim 12.5 h^{-1} \text{Mpc}$) is too small to test if the galaxy luminosity function is a good match to reality, at least the distribution of morphological types is quite realistic: two thirds of the galaxies in this simulation have B -band D/T ratio greater than 0.5. In parallel with this ‘DMG’ simulation, we analysed its dark matter only counterpart, ‘DMO’.

The main effect of the galaxy forming in the halo is to spin up the inner parts, increasing the angular momentum of the region within $0.1 R_{\text{vir}}$ by about 50 per cent in the median. The increase becomes smaller with radius and, at R_{vir} , individual haloes simulated with and without baryons are indistinguishable from each other. We have shown that this increase is due, at least in part, to transfer of angular momentum by gravitational tidal torques, from the baryons to the dark matter. Overall the baryons do not conserve

their own angular momentum, particularly, as Zavala et al. (2008) have shown, in spheroid-dominated galaxies.

The process of galaxy formation also affects the coherence of the halo angular momentum. Although, again, there is a large amount of variation amongst haloes, the median of the distribution indicates that, overall, the formation of the galaxy helps to align the inner angular momentum vector with that of the halo as a whole. An analogous result for the halo shape was obtained by Bailin et al. (2005).

Although the galaxy helps to align the inner halo with its outer parts, the distribution of alignments is still very broad, ranging from perfect alignment to misalignment of $\sim 120^\circ$. Furthermore, because of the large scatter and the relatively small size of our galaxy sample, the distribution of alignments between the inner and outer haloes (which is similar to that between the galaxy and the halo in the DMG simulation) is statistically similar in our simulations with and without baryons. We are also unable to distinguish between the alignment distributions of bulge-dominated and disc-dominated galaxies.

The minor axis of a galaxy is generally oriented parallel to the direction of its angular momentum, even in slowly rotating, bulge-dominated galaxies. In simple models of galaxy formation, we expect the angular momentum of the galaxy to be aligned with the shape of its host halo. This expectation is the basis of the commonly used method of stacking many galaxies together in order to amplify the weak gravitational lensing signal produced by an individual galaxy, in order to estimate, for example, halo ellipticity. In principle, the shape of haloes provides a clean test of the Λ CDM model since, as many studies have shown, CDM haloes are generally strongly aspherical (e.g. Frenk et al. 1988; Jing & Suto 2002; Bett et al. 2007; Hayashi et al. 2007). In particular, such a test could distinguish Λ CDM from modified gravity theories, like TeVeS, in which the potential at large distances from a galaxy should be spherical. Unfortunately, our analysis suggests that this test will not work in practice. The galaxies are sufficiently misaligned with the shape of their haloes that the stacking procedure will wash out any ellipticity signal in the projected mass distribution. Further work with larger galaxy samples is required to ascertain if this problem could be mitigated by a careful sample selection (e.g. according to morphology or luminosity).

In summary, a consistent picture for the spin and shape of Λ CDM haloes seems to be emerging from this and other studies. In the absence of baryons, dark matter haloes are triaxial, with a preference for prolateness, and become increasingly prolate towards their centres (e.g. Bett et al. 2007; Hayashi et al. 2007). They have very little coherent rotation but they rotate with approximately constant rotational rather than angular velocity (Figs. 2 and 11). A galaxy forming in the halo produces several effects: the halo becomes more spherical overall, with a tendency towards oblateness (Kazantzidis et al. 2004; Bailin et al. 2005; Gustafsson et al. 2006; Berentzen & Shlosman 2006; Abadi et al. 2009); its inner parts are spun up (Figs. 11 and 12), and become better aligned with the outer parts, at R_{vir} (Figs. 14 and 15). Thus, as they cool and collapse, the baryons carry information about the outer halo orientation (at R_{vir}) into the inner halo (Figs. 16, 17 and 18; Bailin et al. 2005, Gustafsson et al. 2006). Nevertheless, the alignment of the galaxy with the outer halo is weak and, furthermore, there is very large halo-to-halo scatter in the size of the angular momentum and in the alignment between the inner and outer halo. Half of the galaxies have their minor axes inclined by more than 45° relative to that of their halo and about 10 per cent of the galaxies lie within 30° of the plane perpendicular to the halo major axis (Figs. 20 and 21).

Although we have studied the largest set of simulated haloes and galaxies available to date, two limitations of our galaxy sample should be kept in mind. Firstly, even this sample is too small to search for trends with morphology or formation history. Secondly, the baryonic physics modelled in the simulations are extremely uncertain. Even though our simulated galaxy population appears realistic, at least in so far as the presence of different morphological types is concerned, it includes only one of many possible treatments of these fundamental, but complex, processes.

ACKNOWLEDGEMENTS

PB acknowledges a PhD studentship from the Science & Technology Facilities Council (STFC), and the support of the Helmholtz Alliance HA-101 ‘Physics at the Terascale’. VRE is a Royal Society University Research Fellow. CSF is a Royal Society Wolfson Research Merit Award holder. TO acknowledges support from the FIRST project based on Grants-in-Aid for Specially Promoted Research by MEXT (16002003), Grant-in-Aid for Scientific Research (S) by JSPS (20224002), and an STFC Rolling Grant. The simulations and analyses used in this paper were carried out as part of the programme of the Virgo Consortium on the Regatta supercomputer of the Computing Centre of the Max-Planck-Society in Garching, and the Cosmology Machine supercomputer at the Institute for Computational Cosmology, Durham.

REFERENCES

- Abadi M. G., Navarro J. F., Fardal M., Babul A., Steinmetz M., 2009, preprint (arXiv:0902.2477)
- Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003a, *ApJ*, 591, 499
- Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003b, *ApJ*, 591, 499
- Agustsson I., Brainerd T. G., 2006, *ApJ*, 650, 550
- Allgood B., Flores R. A., Primack J. R., Kravtsov A. V., Wechsler R. H., Faltenbacher A., Bullock J. S., 2006, *MNRAS*, 367, 1781
- Avila-Reese V., Colín P., Gottlöber S., Firmani C., Maulbetsch C., 2005, *ApJ*, 634, 51
- Azzaro M., Patiri S. G., Prada F., Zentner A. R., 2007, *MNRAS*, 376, L43
- Bailin J., Kawata D., Gibson B. K., et al., 2005, *ApJ*, 627, L17
- Bailin J., Steinmetz M., 2004, *ApJ*, 616, 27
- Bailin J., Steinmetz M., 2005, *ApJ*, 627, 647
- Barnes J., Efstathiou G., 1987, *ApJ*, 319, 575
- Barnes J., White S. D. M., 1984, *MNRAS*, 211, 753
- Basilakos S., Plionis M., Maddox S. J., 2000, *MNRAS*, 316, 779
- Baugh C. M., 2006, *Reports of Progress in Physics*, 69, 3101
- Baugh C. M., Lacey C. G., Frenk C. S., Granato G. L., Silva L., Bressan A., Benson A. J., Cole S., 2005, *MNRAS*, 356, 1191
- Bekenstein J. D., 2004, *Phys. Rev. D*, 70, 083509
- Benson A. J., Džanović D., Frenk C. S., Sharples R., 2007, *MNRAS*, 379, 841
- Berentzen I., Shlosman I., 2006, *ApJ*, 648, 807
- Bett P., Eke V., Frenk C. S., Jenkins A., Helly J., Navarro J., 2007, *MNRAS*, 376, 215
- Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*. Princeton University Press, Princeton, NJ USA
- Blumenthal G. R., Faber S. M., Flores R., Primack J. R., 1986, *ApJ*, 301, 27
- Brainerd T. G., 2005, *ApJ*, 628, L101
- Brook C. B., Kawata D., Martel H., Gibson B. K., Bailin J., 2006, *ApJ*, 639, 126
- Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
- Bullock J. S., Dekel A., Kolatt T. S., et al., 2001, *ApJ*, 555, 240
- Bullock J. S., Kolatt T. S., Sigad Y., et al., 2001, *MNRAS*, 321, 559
- Ceverino D., Klypin A., 2008, in Bureau M., Athanassoula E., Barbu B., eds, *IAU Symposium Vol. 245 of IAU Symposium, The role of stellar feedback in the formation of galactic disks and bulges in a Λ CDM Universe*. pp 33–34
- Chen D. N., Jing Y. P., 2002, *MNRAS*, 336, 55
- Chen D. N., Jing Y. P., Yoshikawa K., 2003, *ApJ*, 597, 35
- Cole S., Lacey C., 1996, *MNRAS*, 281, 716
- Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, 319, 168
- Croft R. A. C., Di Matteo T., Springel V., Hernquist L., 2009, *MNRAS*, 400, 43
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- de Jong R. S., Lacey C., 2000, *ApJ*, 545, 781
- Diemand J., Kuhlen M., Madau P., 2007, *ApJ*, 667, 859
- D’Onghia E., Navarro J. F., 2007, *MNRAS*, 380, L58
- Dubinski J., 1992, *ApJ*, 401, 441
- Dubinski J., Carlberg R. G., 1991, *ApJ*, 378, 496
- Dutton A. A., van den Bosch F. C., Dekel A., Courteau S., 2007, *ApJ*, 654, 27
- Eke V. R., Cole S., Frenk C. S., 1996, *MNRAS*, 282, 263
- Evans A. K. D., Bridle S., 2008, *ApJ*, accepted (arXiv:0806.2723)
- Fall S. M., Efstathiou G., 1980, *MNRAS*, 193, 189
- Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, *ApJ*, 327, 507
- Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1985, *Nature*, 317, 595
- Gao L., Navarro J. F., Cole S., Frenk C. S., White S. D. M., Springel V., Jenkins A., Neto A. F., 2008, *MNRAS*, 387, 536
- Gibson B. K., Courty S., Sánchez-Blázquez P., Teyssier R., House E. L., Brook C. B., Kawata D., 2009, in Andersen J., Bland-Hawthorn J., Nordström B., eds, *IAU Symposium Vol. 254 of IAU Symposium, Hydrodynamical Adaptive Mesh Refinement Simulations of Disk Galaxies*. pp 445–452
- Gnedin O. Y., Kravtsov A. V., Klypin A. A., Nagai D., 2004, *ApJ*, 616, 16
- González J. E., Lacey C. G., Baugh C. M., Frenk C. S., Benson A. J., 2009, *MNRAS*, 397, 1254
- Governato F., Mayer L., Wadsley J., Gardner J. P., Willman B., Hayashi E., Quinn T., Stadel J., Lake G., 2004, *ApJ*, 607, 688
- Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, *MNRAS*, 374, 1479
- Gustafsson M., Fairbairn M., Sommer-Larsen J., 2006, *Phys. Rev. D*, 74, 123522
- Haardt F., Madau P., 1996, *ApJ*, 461, 20
- Hahn O., Porciani C., Carollo C. M., Dekel A., 2007, *MNRAS*, 375, 489
- Hayashi E., Navarro J. F., Springel V., 2007, *MNRAS*, 377, 50
- Heller C. H., Shlosman I., Athanassoula E., 2007, *ApJ*, 671, 226
- Hoekstra H., Jain B., 2008, *Annual Review of Nuclear and Particle Science*, 58, 99
- Hoekstra H., Yee H. K. C., Gladders M. D., 2004, *ApJ*, 606, 67
- Jing Y. P., Suto Y., 2002, *ApJ*, 574, 538
- Kang X., van den Bosch F. C., Yang X., Mao S., Mo H. J., Li C., Jing Y. P., 2007, *MNRAS*, 378, 1531

- Kasun S. F., Evrard A. E., 2005, *ApJ*, 629, 781
- Kaufmann T., Mayer L., Wadsley J., Stadel J., Moore B., 2007, *MNRAS*, 375, 53
- Kazantzidis S., Kravtsov A. V., Zentner A. R., Allgood B., Nagai D., Moore B., 2004, *ApJ*, 611, L73
- Knebe A., Kravtsov A. V., Gottlöber S., Klypin A. A., 2000, *MNRAS*, 317, 630
- Lacey C., Cole S., 1994, *MNRAS*, 271, 676
- Lee J., Suto Y., 2004, *ApJ*, 601, 599
- Libeskind N. I., Cole S., Frenk C. S., Okamoto T., Jenkins A., 2007, *MNRAS*, 374, 16
- Macciò A. V., Dutton A. A., van den Bosch F. C., 2008, *MNRAS*, 391, 1940
- Macciò A. V., Dutton A. A., van den Bosch F. C., Moore B., Potter D., Stadel J., 2007, *MNRAS*, 378, 55
- Maller A. H., Dekel A., Somerville R., 2002, *MNRAS*, 329, 423
- Mandelbaum R., Hirata C. M., Broderick T., Seljak U., Brinkmann J., 2006, *MNRAS*, 370, 1008
- Mandelbaum R., Hirata C. M., Seljak U., et al., 2005, *MNRAS*, 361, 1287
- Mayer L., Governato F., Kaufmann T., 2008, *Advanced Science Letters*, 1, 7
- Mo H. J., Mao S., White S. D. M., 1998, *MNRAS*, 295, 319
- Mortlock D. J., Turner E. L., 2001, *MNRAS*, 327, 557
- Natarajan P., Refregier A., 2000, *ApJ*, 538, L113
- Navarro J. F., Benz W., 1991, *ApJ*, 380, 320
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Neto A. F., Gao L., Bett P., et al., 2007, *MNRAS*, 381, 1450
- Okamoto T., 2008, in Frebel A., Maund J. R., Shen J., Siegel M. H., eds, *New Horizons in Astronomy Vol. 393 of Astronomical Society of the Pacific Conference Series, Galaxy Formation*. pp 111–125
- Okamoto T., Eke V. R., Frenk C. S., Jenkins A., 2005, *MNRAS*, 363, 1299
- Okamoto T., Habe A., 1999, *ApJ*, 516, 591
- Parker L. C., Hoekstra H., Hudson M. J., van Waerbeke L., Mellier Y., 2007, *ApJ*, 669, 21
- Parry O. H., Eke V. R., Frenk C. S., 2009, *MNRAS*, 396, 1972
- Pearce F. R., Jenkins A., Frenk C. S., White S. D. M., Thomas P. A., Couchman H. M. P., Peacock J. A., Efstathiou G., 2001, *MNRAS*, 326, 649
- Plionis M., Barrow J. D., Frenk C. S., 1991, *MNRAS*, 249, 662
- Porciani C., Dekel A., Hoffman Y., 2002a, *MNRAS*, 332, 325
- Porciani C., Dekel A., Hoffman Y., 2002b, *MNRAS*, 332, 339
- Reblinsky K., 2000, *A&A*, 364, 377
- Robertson B., Yoshida N., Springel V., Hernquist L., 2004, *ApJ*, 606, 32
- Robertson B. E., Kravtsov A. V., 2008, *ApJ*, 680, 1083
- Romano-Díaz E., Shlosman I., Heller C., Hoffman Y., 2009, *ApJ*, 702, 1250
- Sales L., Lambas D. G., 2004, *MNRAS*, 348, 1236
- Sales L., Lambas D. G., 2009, *MNRAS*, 395, 1184
- Salpeter E. E., 1955, *ApJ*, 121, 161
- Scannapieco C., White S. D. M., Springel V., Tissera P. B., 2009, *MNRAS*, 396, 696
- Sereno M., De Filippis E., Longo G., Bautz M. W., 2006, *ApJ*, 645, 170
- Sharma S., Steinmetz M., 2005, *ApJ*, 628, 21
- Shaw L. D., Weller J., Ostriker J. P., Bode P., 2006, *ApJ*, 646, 815
- Sommer-Larsen J., Götz M., Portinari L., 2003, *ApJ*, 596, 47
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., Wang J., Vogelsberger M., Ludlow A., Jenkins A., Helmi A., Navarro J. F., Frenk C. S., White S. D. M., 2008, *MNRAS*, 391, 1685
- Springel V., White S. D. M., Hernquist L., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, *Dark Matter in Galaxies Vol. 220 of IAU Symposium, The shapes of simulated dark matter halos*. pp 421–429
- Springel V., White S. D. M., Jenkins A., et al., 2005, *Nature*, 435, 629
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V., Yoshida N., White S. D. M., 2001, *New Astronomy*, 6, 79
- Sutherland R. S., Dopita M. A., 1993, *ApJS*, 88, 253
- Tasker E. J., Bryan G. L., 2006, *ApJ*, 641, 878
- van den Bosch F. C., Abel T., Croft R. A. C., Hernquist L., White S. D. M., 2002, *ApJ*, 576, 21
- van den Bosch F. C., Abel T., Hernquist L., 2003, *MNRAS*, 346, 177
- Vitvitska M., Klypin A. A., Kravtsov A. V., Wechsler R. H., Primack J. R., Bullock J. S., 2002, *ApJ*, 581, 799
- Wang Y., Yang X., Mo H. J., Li C., van den Bosch F. C., Fan Z., Chen X., 2008, *MNRAS*, 385, 1511
- Warren M. S., Quinn P. J., Salmon J. K., Zurek W. H., 1992, *ApJ*, 399, 405
- White S. D. M., Frenk C. S., 1991, *ApJ*, 379, 52
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Yang X., van den Bosch F. C., Mo H. J., Mao S., Kang X., Weinmann S. M., Guo Y., Jing Y. P., 2006, *MNRAS*, 369, 1293
- Yoshida N., Abel T., Hernquist L., Sugiyama N., 2003, *ApJ*, 592, 645
- Zaroubi S., Squires G., de Gasperis G., Evrard A. E., Hoffman Y., Silk J., 2001, *ApJ*, 561, 600
- Zavala J., Okamoto T., Frenk C. S., 2008, *MNRAS*, 387, 364

APPENDIX A: ORIENTATION RESOLUTION TESTS

In Sections 2.5.2 and 3.4.2 we studied the orientation of halo (or galaxy) angular momentum vectors. These are formed from a vector sum of the contributions from the objects’ constituent particles. Because of the vector sum, the contribution from many particles can be much smaller than that of a single particle. If the summed vector’s magnitude is particularly small, then the individual particles’ vectors must have been largely in opposite directions, and mostly cancelled out. In that case, the inclusion of very few particles with parallel vectors, or even a single particle, can completely dominate the result. Clearly, any property that is dominated by discreteness effects is not reliable. To ensure our results are robust, we have performed extensive Monte Carlo tests using the halo catalogues from the hMS, DMO and DMG simulations. (This problem has been tackled in the past by Bullock et al. (2001), Bailin & Steinmetz (2004), and Avila-Reese et al. (2005). We revisit it here to ensure that we can retain as many haloes as possible, while rejecting those that are unreliable.)

For each halo, we perform 5000 bootstrap resamplings of both the particles from the halo as a whole and also from just the halo inner region, independently. The cosine of the angle between the original \mathbf{j}_{tot} (or $\mathbf{j}_{\text{inner}}$) and the bootstrap resampled \mathbf{j}_{tot} (or $\mathbf{j}_{\text{inner}}$) is computed. We take the median of the 5000 samples for each halo and plot it against the magnitude of the original vector, rescaled in

such a manner that any systematic trends are removed. For this, we use the dimensionless quantity $\tilde{j}(\leq r) := j(\leq r)/j_0(r)$, where $j_0(r)$ is the specific angular momentum of a (hypothetical) test particle in a circular orbit at radius r . Since the circular velocity $v_c(r) = \sqrt{GM(\leq r)r^{-1}}$, we have:

$$\tilde{j}(\leq r) = \frac{j(\leq r)}{\sqrt{GM(\leq r)r}} \quad (\text{A1})$$

We have confirmed that this quantity does not vary systematically with either mass or radius. (This \tilde{j} is similar to, but not the same as, the scaled angular momentum shown in Fig. 9.)

We wish to find a limiting value of $\tilde{j} =: \tilde{j}_{\text{lim}}$ such that some given percentage of haloes with $\tilde{j} \geq \tilde{j}_{\text{lim}}$ have their median bootstrapped vectors aligned to within a given angle. Specifically, we require that 99.5 per cent of the selected haloes should have $\theta_{\text{med}} \leq 15^\circ$, where θ is the angle between the actual halo \mathbf{j} and the bootstrap resampled vector.

In practice, the limits obtained from this process, and the severity of the cut on the resulting halo population, depend strongly on the other selection criteria used. We always restrict attention to ‘virialised’ haloes ($Q_{\text{lim}} = 0.5$), but the minimum number of particles in the given region (total or inner halo) has a strong impact. To balance the competing demands of a well resolved inner region and spin orientations that are robust to discreteness effects, we decided to move the ‘inner’ radius outwards. While for the spin magnitude analysis we adopted $r_{\text{inner}} = 0.1R_{\text{vir}}$, for the orientation profiles we chose to move this two bins outwards to $r_{\text{inner}} = 10^{-0.6}R_{\text{vir}} \approx 0.25R_{\text{vir}}$. We also take 1000 as the minimum number of particles in the given region (total or inner), rather than the 300 we use for the analysis of the angular momentum magnitude. With these choices we were left with a large enough sample of haloes to be statistically viable once the cuts in N_p and \tilde{j} had been applied.

Fig. A1 shows the bootstrap results for the angular momentum of dark matter within R_{vir} . The median of the bootstrapped angles from each halo are plotted. There is a clear increase in scatter for haloes with low \tilde{j} (that is, low j compared to that of a typical particle). The same trend is present for haloes from the hMS, DMO and DMG simulations. Our requirement that most haloes should have a median bootstrap angle within 15° translates into a selection criterion of approximately $\log_{10} \tilde{j} \geq -1.5$. Fig. A2 shows the results of applying this method to the inner regions of the haloes only. Although the scatter behaves slightly differently, the value of \tilde{j}_{lim} is very similar. It is worth noting that, as long as they are selected with the same criteria, the inclusion of the DMO and DMG haloes does not affect the value of \tilde{j}_{lim} significantly. Including the stellar components of galaxies does not affect the results greatly either, since the galaxies generally have higher angular momentum anyway.

Thus, the halo selection criteria we adopt for our analysis of angular momentum orientation are:

$$\begin{aligned} Q &\leq 0.5 \\ N_p &\geq 1000 \text{ (inner or total)} \\ \log_{10} \tilde{j} &\geq \begin{cases} -1.44 & \text{(total)} \\ -1.51 & \text{(inner)} \end{cases} \end{aligned}$$

We do not select according to the \tilde{j} of the galaxies.

A similar problem arises when considering the orientation of the dark matter mass distribution, i.e. the halo shape given by the mutually orthogonal vectors \mathbf{a} , \mathbf{b} and \mathbf{c} . Regarding the axis ratio $s = c/a$ as a measure of how close the halo is to spherical, it is clear that if s is very close to unity, then the three axes become de-




Figure A1. Bootstrap resampling results for the angular momentum of the dark matter within R_{vir} . Each point is the median of the 5000 angles between the \mathbf{j}_{tot} of a halo and that of the bootstrap resamplings of that halo. This angle is plotted against the halo’s scaled angular momentum \tilde{j} , so that the different simulations can be compared together. The heavy magenta contour joins the series of lower limits on \tilde{j} such that 99.5 per cent of the haloes with $\tilde{j} \geq \tilde{j}_{\text{lim}}$ are better aligned in the median than that angle. The haloes have been selected to be ‘virialised’ ($Q \leq 0.5$) and well resolved (at least 1000 dark matter particles), and have been taken from the hMS (black dots), DMO (cyan squares) and DMG (red rings) simulations.




Figure A2. As for Fig. A1, but for the angular momentum of the dark matter within $\approx 0.25R_{\text{vir}}$; the N_p selection criterion therefore applies to this inner region. hMS haloes are marked as black dots, DMO haloes as blue crosses, and DMG haloes as red asterisks. Also included are the angular momenta of the stellar components of the central galaxies from the DMG haloes (green rings). To be included, the DMG haloes/galaxies need also to satisfy $N_p \geq 1000$ for the star particles. Inclusion of the galaxies does not affect the resulting limit on \tilde{j} .




Figure A3. The angle between the measured and bootstrapped halo major axes plotted against the minor-to-major axis ratio, s , using similar principles to those in Fig A1. Due to the greater scatter in the angles in this case, we plot two contours. One (thin) is such that 99.5 per cent of the haloes that have a larger value of s are better aligned than the given angle; the other (thick) corresponds to the 95.0 per cent level. The larger scatter results in much more restrictive cuts.

generate and the halo orientation is undetermined. Thus, just as for angular momentum, we use bootstrap resampling to find a limiting value of s (s_{max}) such that an acceptable fraction of the haloes with $s \leq s_{\text{max}}$ have sufficiently reliable directions (Fig. A3). We find that the shape axes scatter more under bootstrapping than the angular momentum vectors. If we want to retain a significantly large sample we have to be less demanding on the accuracy of their orientations. We reduce the percentile limit to be such that 95 per cent of haloes with $s \leq s_{\text{max}}$ have their median bootstrap angle within 15° . This yields a limit of $s_{\text{max}} = 0.81$ which we apply as an additional selection criterion in Figs. 6, 19, and 20.